noise from images having textures or very detailed areas or prominent edges. Smoothing by low-pass and similar linear filtering techniques, whereby variations in pixel intensities within a given area are averaged out, is not effective since it often results in reconstructed, or repaired, scratch noise areas being blurred and loss of detail in the repaired scratch area.

With median filtering techniques, a given pixel of a subject image is replaced with the median value of the surrounding pixels. Which pixels comprise the "surrounding" pixels of a pixel at issue is typically defined by specifying the limits of a window enclosing that pixel. For example a window of 10 pixels by 10 pixels (10×10) in which the pixel at issue is located could be used to define the "surrounding" pixels. The surrounding pixels can also be referred to as the "neighborhood" of a given pixel. The size of the window or neighborhood can be varied in accordance with the needs at hand. In applying median filtering to scratch removal where the location of the scratch is known, only the scratch pixels will be replaced by the median value of the surrounding pixels.

Median and similar non-linear filtering techniques work relatively well where the scratch area is limited to a relatively small group of isolated pixels. Typically scratches, and particularly wire noise, consist of many contiguous pixels. As a result the median value of surrounding pixels is typically the value of the scratch area itself. If a small window, or neighborhood, is chosen, scratches are typically not repaired, since the median value is typically the value of the scratch area itself. In other words, the scratch pixels values are replaced with the value of the scratch pixels. Where a larger window is chosen so as to provide a greater number of surrounding pixels, blurring of the image typically results.

In techniques based on statistical texture synthesis, an area of the damaged image at issue is analyzed and a statistical representation of the image area is developed. This statistical representation is then used to generate a synthesized texture that is similar to the analyzed area. Scratch/wire noise pixels can then be replaced with the synthesized texture. Statistical methods are based on the assumption that the relationship between pixels in a neighborhood has a statistical description which can be discovered by analysis. Thus a given pixel's characteristics are supposed to depend statistically on some or all of it's nearby pixels. Where there is a relationship between pixels which are relatively far apart in distance, the neighborhood necessary for analysis of those pixels must be relatively large. An example would be a brick wall with large bricks or a wide wire mesh photographed close-up. Typically in these methods the computational processing required increases as the number of pixels within the neighborhood on which a given pixel depends is increased. As a result these methods are not useful for textures having large features, such as, for example, the texture of an image of a typical brick wall.

Many popular image processing programs, such as for example, Adobe Photoshop™, provide a tool for copying pixels from one area of the image to another area. With these types of tools, the area of the image to be copied is typically selected, or defined, via use of a mouse, or other pointer/control device, which allows a user to manipulate a cursor or "brush" on the display screen which displays the image at issue. Using the mouse device, a source area from which pixel data is to be read, or transferred, is defined. Then a destination area to which the pixel data from the source area is to be transferred is defined. "Brush strokes", or movement of the cursor across the destination area, can be accomplished using the mouse. With each "brush stroke" pixels from the source area are copied to the destination area. This method works well when the texture, or pattern, depicted in the image at issue is random or regularly occurring in nature and has relatively small features. Small features might include, for example, spots on a piece of fruit, sand, asphalt of a road, or fabric weave. Where there are, for example, random or regularly occurring details, or patterns, depicted in the image in the form of curved or straight lines then alignment of the edges of the repaired region with that of the surrounding region is very time consuming. Repairing scratches or wires using rubber stamp painting is a task requiring lot of time and skill for many images. Since the human visual system is very sensitive to noticing problems with edges, mismatched edges are relatively easy to pick out by the human eye.

Another technique for scratch removal has been to "paint" over the pixels of the scratch by manually matching the scratch pixels with the surrounding area. This method is very time consuming and the quality of the results obtained depends heavily on the skill of the operator.

Techniques of scratch noise removal based on projection methods in the context of signal restoration have been previously used. These methods have been used to extend the frequency spectrum of band-limited signals by alternating between signal and transform domains and applying constraints of each domain. The constraints are directed to preservation of known signal values and enforcement of bandlimits. This approach has been generalized through a geometric interpretation which is generally referred to as Projection onto Convex Sets, or POCS. POCS allows the use of any information about the image, such as, for example, the range of luminance values, distinction between defined scratch areas and non-scratch areas, as long as the information can be represented as a closed convex set to be used as a constraint.

In projection based methods an image at issue is projected upon predefined constraints to find another image which closely resembles the image at issue and which is within a defined range of images which satisfy certain pre-defined constraints. The process of finding the closest image satisfying certain pre-defined constraints is known as projecting onto that constraint. Once such an image has been found, this new image, or resultant image, is projected onto the next pre-defined constraint, if any. The constraints represent information that is known about the image. In projection based methods, an image at issue may be considered as a single point in a space which has one or more dimensions. Any point on a flat surface can be described in terms of its X and Y coordinates. Further any point in a 3-dimensional space, for example a room, can be described in terms of X, Y and Z coordinates. Similarly, an image can be described as a single point in high-dimensional space. For example, an image comprised of 200 pixels by 100 pixels may be represented as a single point having 20,000 dimensions (200×100=20,000). Each dimension of the point represents one of the 20,000 pixels which make up the image. The value of each dimension represents the luminance intensity for the given pixel represented by the particular dimension. In typical gray scale imaging, a gray scale of 256 steps is often used for high quality imaging. Thus, the luminance intensity, or value, of a given pixel may range from 0–255, for example.

Through repeated projections onto various constraints an image can eventually be found that satisfies all of the constraints imposed, simultaneously. The only condition is that the constraints should represent something known as